An overview of embedding models of entities and relationships for knowledge base completion

Dat Quoc Nguyen

Department of Computing
Macquarie University, Sydney, Australia
dat.nguyen@students.mq.edu.au

arXiv:1703.08098v2 [cs.CL] 28 Mar 2017

Abstract

Knowledge bases of real-world facts about entities and their relationships are useful resources for a variety of natural language processing tasks. However, because knowledge bases are typically incomplete, it is useful to be able to perform *knowledge base completion* or *link prediction*, i.e., predict whether a relationship not in the knowledge base is likely to be true. This article presents an overview of embedding models of entities and relationships for knowledge base completion.

Keywords: Knowledge base completion, link prediction, embedding model, triple classification, entity prediction.

1 Introduction

Knowledge bases (KBs), such as WordNet (Miller, 1995), YAGO (Suchanek et al., 2007), Freebase (Bollacker et al., 2008) and DBpedia (Lehmann et al., 2015), represent relationships between entities as triples $(\text{head entity}, \text{relation}, \text{tail entity})$. Even very large knowledge bases are still far from complete (Socher et al., 2013; West et al., 2014). *Knowledge base completion* or *link prediction* systems (Nickel et al., 2016a) predict which triples not in a knowledge base are likely to be true (Taskar et al., 2004; Bordes et al., 2011).

Embedding models for KB completion associate entities and/or relations with dense feature vectors or matrices. Such models obtain the state-of-the-art performances (Bordes et al., 2013; Wang et al., 2014; Guu et al., 2015; Nguyen et al., 2016a; Nguyen et al., 2016b) and generalize to large KBs (Krompaß et al., 2015). This article serves as a brief overview of embedding models for KB completion, with up-to-date experimental results on two standard evaluation tasks: i) the entity prediction task—which is also referred to as the link prediction task (Bordes et al., 2013)—and ii) the triple classification task (Socher et al., 2013).

2 Embedding models for KB completion

2.1 General idea

Let $\mathcal{E}$ denote the set of entities and $\mathcal{R}$ the set of relation types. Denote by $\mathcal{G}$ the knowledge base consisting of a set of correct triples (h, r, t), such that $h, t \in \mathcal{E}$ and $r \in \mathcal{R}$. For each triple (h, r, t), the embedding models define a *score function* $f(h, r, t)$ of its implausibility. Their goal is to choose f such that the score $f(h, r, t)$ of a plausible triple (h, r, t) is smaller than the score $f(h', r', t')$ of an implausible triple (h', r', t').

Table 1 summarizes different score functions $f(h, r, t)$ and the optimization algorithms used to estimate model parameters, e.g., vanilla Stochastic Gradient Descent (SGD), SGD+AdaGrad (Duchi et al., 2011), SGD+AdaDelta (Zeiler, 2012) and L-BFGS (Liu and Nocedal, 1989). To learn model parameters (i.e. entity vectors, relation vectors or matrices), the embedding models minimize an objective function. A common objective function is the following margin-based function:

$$\mathcal{L} = \sum_{\substack{(h,r,t)\in\mathcal{G} \\ (h',r,t')\in\mathcal{G}'_{(h,r,t)}}} [\gamma + f(h, r, t) - f(h', r, t')]_+$$

where $[x]_+ = \max(0, x)$, γ is the margin hyper-parameter, and

$$\mathcal{G}'_{(h,r,t)} = \{(h', r, t) \mid h' \in \mathcal{E}, (h', r, t) \notin \mathcal{G}\} \cup \{(h, r, t') \mid t' \in \mathcal{E}, (h, r, t') \notin \mathcal{G}\}$$

is the set of incorrect triples generated by corrupting the correct triple $(h, r, t) \in \mathcal{G}$.

Model	Score function $f(h, r, t)$	**Opt.**
Unstructured	$\|\boldsymbol{v}_h - \boldsymbol{v}_t\|_{\ell_{1/2}}$	SGD
SE	$\|\mathbf{W}_{r,1}\boldsymbol{v}_h - \mathbf{W}_{r,2}\boldsymbol{v}_t\|_{\ell_{1/2}}$; $\mathbf{W}_{r,1}, \mathbf{W}_{r,2} \in \mathbb{R}^{k\times k}$	SGD
SME	$(\mathbf{W}_{1,1}\boldsymbol{v}_h + \mathbf{W}_{1,2}\boldsymbol{v}_r + \mathbf{b}_1)^\top(\mathbf{W}_{2,1}\boldsymbol{v}_t + \mathbf{W}_{2,2}\boldsymbol{v}_r + \mathbf{b}_2)$ $\mathbf{b}_1, \mathbf{b}_2 \in \mathbb{R}^n$; $\mathbf{W}_{1,1}, \mathbf{W}_{1,2}, \mathbf{W}_{2,1}, \mathbf{W}_{2,2} \in \mathbb{R}^{n\times k}$	SGD
TransE	$\|\boldsymbol{v}_h + \boldsymbol{v}_r - \boldsymbol{v}_t\|_{\ell_{1/2}}$; $\boldsymbol{v}_r \in \mathbb{R}^k$	SGD
TransH	$\|(\mathbf{I} - \boldsymbol{r}_p\boldsymbol{r}_p^\top)\boldsymbol{v}_h + \boldsymbol{v}_r - (\mathbf{I} - \boldsymbol{r}_p\boldsymbol{r}_p^\top)\boldsymbol{v}_t\|_{\ell_{1/2}}$ $\boldsymbol{r}_p, \boldsymbol{v}_r \in \mathbb{R}^k$; **I**: Identity matrix size $k \times k$	SGD
TransR	$\|\mathbf{W}_r\boldsymbol{v}_h + \boldsymbol{v}_r - \mathbf{W}_r\boldsymbol{v}_t\|_{\ell_{1/2}}$; $\mathbf{W}_r \in \mathbb{R}^{n\times k}$; $\boldsymbol{v}_r \in \mathbb{R}^n$	SGD
TransD	$\|(\mathbf{I} + \boldsymbol{r}_p\boldsymbol{h}_p^\top)\boldsymbol{v}_h + \boldsymbol{v}_r - (\mathbf{I} + \boldsymbol{r}_p\boldsymbol{t}_p^\top)\boldsymbol{v}_t\|_{\ell_{1/2}}$ $\boldsymbol{r}_p, \boldsymbol{v}_r \in \mathbb{R}^n$; $\boldsymbol{h}_p, \boldsymbol{t}_p \in \mathbb{R}^k$; **I**: Identity matrix size $n \times k$	AdaDelta
lppTransD	$\|(\mathbf{I} + \boldsymbol{r}_{p,1}\boldsymbol{h}_p^\top)\boldsymbol{v}_h + \boldsymbol{v}_r - (\mathbf{I} + \boldsymbol{r}_{p,2}\boldsymbol{t}_p^\top)\boldsymbol{v}_t\|_{\ell_{1/2}}$ $\boldsymbol{r}_{p,1}, \boldsymbol{r}_{p,2}, \boldsymbol{v}_r \in \mathbb{R}^n$; $\boldsymbol{h}_p, \boldsymbol{t}_p \in \mathbb{R}^k$; **I**: Identity matrix size $n \times k$	SGD
TranSparse	$\|\mathbf{W}_r^h(\theta_r^h)\boldsymbol{v}_h + \boldsymbol{v}_r - \mathbf{W}_r^t(\theta_r^t)\boldsymbol{v}_t\|_{\ell_{1/2}}$; $\mathbf{W}_r^h, \mathbf{W}_r^t \in \mathbb{R}^{n\times k}$; $\theta_r^h, \theta_r^t \in \mathbb{R}$; $\boldsymbol{v}_r \in \mathbb{R}^n$	SGD
STransE	$\|\mathbf{W}_{r,1}\boldsymbol{v}_h + \boldsymbol{v}_r - \mathbf{W}_{r,2}\boldsymbol{v}_t\|_{\ell_{1/2}}$; $\mathbf{W}_{r,1}, \mathbf{W}_{r,2} \in \mathbb{R}^{k\times k}$; $\boldsymbol{v}_r \in \mathbb{R}^k$	SGD
DISTMULT	$\boldsymbol{v}_h^\top\mathbf{W}_r\boldsymbol{v}_t$; $\mathbf{W}_r$ is a diagonal matrix $\in \mathbb{R}^{k\times k}$	AdaGrad
NTN	$\boldsymbol{v}_r^\top tanh(\boldsymbol{v}_h^\top\mathbf{M}_r\boldsymbol{v}_t + \mathbf{W}_{r,1}\boldsymbol{v}_h + \mathbf{W}_{r,2}\boldsymbol{v}_t + \mathbf{b}_r)$ $\boldsymbol{v}_r, \mathbf{b}_r \in \mathbb{R}^n$; $\mathbf{M}_r \in \mathbb{R}^{k\times k\times n}$; $\mathbf{W}_{r,1}, \mathbf{W}_{r,2} \in \mathbb{R}^{n\times k}$	L-BFGS
HolE	$sigmoid(\boldsymbol{v}_r^\top(\boldsymbol{v}_h \circ \boldsymbol{v}_t))$; $\boldsymbol{v}_r \in \mathbb{R}^k$, $\circ$ denotes *circular correlation*	AdaGrad
Bilinear-COMP	$\boldsymbol{v}_h^\top\mathbf{W}_{r_1}\mathbf{W}_{r_2}...\mathbf{W}_{r_m}\boldsymbol{v}_t$; $\mathbf{W}_{r_1}, \mathbf{W}_{r_2}, ..., \mathbf{W}_{r_m} \in \mathbb{R}^{k\times k}$	AdaGrad
TransE-COMP	$\|\boldsymbol{v}_h + \boldsymbol{v}_{r_1} + \boldsymbol{v}_{r_2} + ... + \boldsymbol{v}_{r_m} - \boldsymbol{v}_t\|_{\ell_{1/2}}$; $\boldsymbol{v}_{r_1}, \boldsymbol{v}_{r_2}, ..., \boldsymbol{v}_{r_m} \in \mathbb{R}^k$	AdaGrad

Table 1: The score functions $f(h, r, t)$ and the optimization methods (Opt.) of several prominent embedding models for KB completion. In all of these models, the entities h and t are represented by vectors $\boldsymbol{v}_h$ and $\boldsymbol{v}_t \in \mathbb{R}^k$, respectively.

2.2 Specific models

The Unstructured model (Bordes et al., 2012) assumes that the head and tail entity vectors are similar. As the Unstructured model does not take the relationship into account, it cannot distinguish different relation types. The Structured Embedding (SE) model (Bordes et al., 2011) extends the Unstructured model by assuming that the head and tail entities are similar only in a relation-dependent subspace, where each relation is represented by two different matrices. Futhermore, the SME model (Bordes et al., 2012) uses four different matrices to project entity and relation vectors into a subspace. The TransE model (Bordes et al., 2013) is inspired by models such as the Word2Vec Skip-gram model (Mikolov et al., 2013) where relationships between words often correspond to translations in latent feature space.

The TransH model (Wang et al., 2014) associates each relation with a relation-specific hyperplane and uses a projection vector to project entity vectors onto that hyperplane. TransD (Ji et al., 2015) and TransR/CTransR (Lin et al., 2015b) extend the TransH model by using two projection vectors and a matrix to project entity vectors into a relation-specific space, respectively. Similar to TransR, TransR-FT (Feng et al., 2016a) also uses a matrix to project head and tail entity vectors. TEKE_H (Wang and Li, 2016) extends TransH to incorporate rich context information in an external text corpus. lppTransD (Yoon et al., 2016) extends TransD to additionally use two projection vectors for representing each relation. STransE (Nguyen et al., 2016b) and TranSparse (Ji et al., 2016) can be viewed as direct extensions of the TransR model, where head and tail entities are associated with their own projection matrices. Unlike STransE, TranSparse uses adaptive sparse matrices, whose sparse degrees are defined based on the number of entities linked by relations.

DISTMULT (Yang et al., 2015) is based on the Bilinear model (Nickel et al., 2011; Bordes et al., 2012; Jenatton et al., 2012) where each relation is represented by a diagonal rather than a full matrix. The neural tensor network (NTN) model (Socher et al., 2013) uses a bilinear tensor operator to represent each relation while ER-MLP (Dong et al., 2014) and ProjE (Shi and Weninger, 2017) could be viewed as simplified versions of NTN. Such quadratic forms are also used to model entities and relations in KG2E (He et al., 2015), TransG (Xiao et al., 2016), ComplEx (Trouillon et al., 2016), TATEC (García-Durán et al., 2016) and RSTE (Tay et al., 2017). In addition, HolE (Nickel et al., 2016b) uses circular correlation—a compositional operator—which could be interpreted as a compression of the tensor product.

Recent research has shown that relation paths between entities in KBs provide richer context information and improve the performance of embedding models for KB completion (Luo et al., 2015; Lin et al., 2015a; Liang and Forbus, 2015; García-Durán et al., 2015; Guu et al., 2015; Toutanova et al., 2016; Nguyen et al., 2016a).

Luo et al. (2015) constructed relation paths between entities and viewing entities and relations in the path as pseudo-words, and then applied Word2Vec algorithms (Mikolov et al., 2013) to produce pre-trained vectors for these pseudo-words. Luo et al. (2015) showed that using these pre-trained vectors for initialization helps to improve the performance of the TransE (Bordes et al., 2013), SME (Bordes et al., 2012) and SE (Bordes et al., 2011). In addition, Liang and Forbus (2015) used the implausibility score produced by SME to compute the weights of relation paths.

Furthermore, RTransE (García-Durán et al., 2015), PTransE (Lin et al., 2015a) and TransE-COMP (Guu et al., 2015) are extensions of the TransE model. These models similarly represent a relation path by a vector which is the sum of the vectors of all relations in the path, whereas in the Bilinear-COMP model (Guu et al., 2015) and the PRUNED-PATHS model (Toutanova et al., 2016), each relation is a matrix and so it represents the relation path by matrix multiplication. The neighborhood mixture model TransE-NMM (Nguyen et al., 2016a) can be also viewed as a three-relation path model as it takes into account the neighborhood entity and relation information of both head and tail entities in each triple.

Other relation path-based models: The Path Ranking Algorithm (PRA) (Lao and Cohen, 2010) is a random walk inference technique which was proposed to predict a new relationship between two entities in KBs. Lao et al. (2011) used PRA to estimate the probability of an unseen triple as a combination of weighted random walks that follow different paths linking the head entity and tail entity in the KB. Gardner et al. (2014) made use of an external text corpus to increase the connectivity of the KB used as the input to PRA. Gardner and Mitchell (2015) improved PRA by proposing a subgraph feature extraction technique to make the generation of random walks in KBs more efficient and expressive, while Wang et al. (2016) extended PRA to couple the path ranking of multiple relations. PRA can also be used in conjunction with first-order formula in the discriminative Gaifman model (Niepert, 2016). In addition, Neelakantan et al. (2015) used a recurrent neural network to learn vector representations of PRA-style relation paths between entities in the KB. Other random-walk based learning algorithms for KB completion can be also found in Feng et al. (2016b), Liu et al. (2016) and Wei et al. (2016).

See other methods for learning from KBs and multi-relational data in Nickel et al. (2016a).

3 Evaluation tasks

Two standard tasks are proposed to evaluate embedding models for KB completion including: the entity prediction task, i.e. link prediction (Bordes et al., 2013), and the triple classification task (Socher et al., 2013).

Dataset	#E	#R	#Train	#Valid	#Test
WN18	40,943	18	141,442	5,000	5,000
FB15k	14,951	1,345	483,142	50,000	59,071
WN11	38,696	11	112,581	2,609	10,544
FB13	75,043	13	316,232	5,908	23,733

Table 2: Statistics of the experimental datasets. #E is the number of entities, #R is the number of relation types, and #Train, #Valid and #Test are the numbers of correct triples in the training, validation and test sets, respectively. In both WN11 and FB13, each validation and test set also contains the same number of incorrect triples as the number of correct triples.

Commonly, the WN18 and FB15k datasets (Bordes et al., 2013) are used for entity prediction evaluation, while the WN11 and FB13 datasets

Method	Raw						Filtered					
	WN18			FB15k			WN18			FB15k		
	MR	H@10	MRR	MR	H@10	MRR	MR	H@10	MRR	MR	H@10	MRR
SE (Bordes et al., 2011)	1011	68.5	-	273	28.8	-	985	80.5	-	162	39.8	-
Unstructured (Bordes et al., 2012)	315	35.3	-	1074	4.5	-	304	38.2	-	979	6.3	-
SME (Bordes et al., 2012)	545	65.1	-	274	30.7	-	533	74.1	-	154	40.8	-
TransH (Wang et al., 2014)	401	73.0	-	212	45.7	-	303	86.7	-	87	64.4	-
TransR (Lin et al., 2015b)	238	79.8	-	198	48.2	-	225	92.0	-	77	68.7	-
CTransR (Lin et al., 2015b)	231	79.4	-	199	48.4	-	218	92.3	-	75	70.2	-
KG2E (He et al., 2015)	342	80.2	-	**174**	48.9	-	331	92.8	-	59	74.0	-
TransD (Ji et al., 2015)	224	79.6	-	194	53.4	-	212	92.2	-	91	77.3	-
lppTransD (Yoon et al., 2016)	283	80.5	-	195	53.0	-	270	94.3	-	78	78.7	-
TranSparse (Ji et al., 2016)	223	80.1	-	187	**53.5**	-	211	93.2	-	82	79.5	-
TATEC (García-Durán et al., 2016)	-	-	-	-	-	-	-	-	-	**58**	76.7	-
NTN (Socher et al., 2013)	-	-	-	-	-	-	-	66.1	0.53	-	41.4	0.25
DISTMULT (Yang et al., 2015)	-	-	-	-	-	-	-	94.2	0.83	-	57.7	0.35
ComplEx (Trouillon et al., 2016)	-	-	0.587	-	-	0.242	-	94.7	**0.941**	-	**84.0**	**0.692**
HolE (Nickel et al., 2016b)	-	-	**0.616**	-	-	0.232	-	**94.9**	0.938	-	73.9	0.524
RESCAL (Nickel et al., 2011) [*]	-	-	0.603	-	-	0.189	-	92.8	0.890	-	58.7	0.354
TransE (Bordes et al., 2013) [*]	-	-	0.351	-	-	0.222	-	94.3	0.495	-	74.9	0.463
STransE (Nguyen et al., 2016b)	**217**	**80.9**	0.469	219	51.6	**0.252**	**206**	93.4	0.657	69	79.7	0.543
RTransE (García-Durán et al., 2015)	-	-	-	-	-	-	-	-	-	**50**	76.2	-
PTransE (Lin et al., 2015a)	-	-	-	207	51.4	-	-	-	-	58	84.6	-
GAKE (Feng et al., 2016b)	-	-	-	228	44.5	-	-	-	-	119	64.8	-
Gaifman (Niepert, 2016)	-	-	-	-	-	-	352	93.9	-	75	84.2	-
Hiri (Liu et al., 2016)	-	-	-	-	-	-	-	90.8	0.691	-	70.3	0.603
NLFeat (Toutanova and Chen, 2015)	-	-	-	-	-	-	-	**94.3**	**0.940**	-	**87.0**	**0.822**
TEKE_H (Wang and Li, 2016)	**127**	80.3	-	212	51.2	-	**114**	92.9	-	108	73.0	-
SSP (Xiao et al., 2017)	168	**81.2**	-	**163**	**57.2**	-	156	93.2	-	82	79.0	-

Table 3: Entity prediction results. MR and H@10 denote evaluation metrics of mean rank and Hits@10 (in %), respectively. "NLFeat" abbreviates Node+LinkFeat. The results for NTN (Socher et al., 2013) listed in this table are taken from Yang et al. (2015) since NTN was originally evaluated only for triple classification. [*]: Results from the implementation of Nickel et al. (2016b) because these results are higher than those previously published in Bordes et al. (2013).

(Socher et al., 2013) are used for triple classification. WN18 and WN11 are derived from the large lexical KB WordNet (Miller, 1995). FB15k and FB13 are derived from the large real-world fact KB FreeBase (Bollacker et al., 2008). Information about these datasets is given in Table 2.

3.1 Entity prediction

3.1.1 Task description

The entity prediction task, i.e. link prediction (Bordes et al., 2013), predicts the head or the tail entity given the relation type and the other entity, i.e. predicting h given $(?, r, t)$ or predicting t given $(h, r, ?)$ where ? denotes the missing element. The results are evaluated using a ranking induced by the function $f(h, r, t)$ on test triples.

Each correct test triple (h, r, t) is corrupted by replacing either its head or tail entity by each of the possible entities in turn, and then these candidates are ranked in ascending order of their implausibility score. This is called as the "Raw" setting protocol. Furthermore, the "Filtered" setting protocol, described in Bordes et al. (2013), filters out before ranking any corrupted triples that appear in the KB. Ranking a corrupted triple appearing in the KB (i.e. a correct triple) higher than the original test triple is also correct, but is penalized by the "Raw" score, thus the "Filtered" setting provides a clearer view on the ranking performance.

In addition to the mean rank and the Hits@10 (i.e., the proportion of test triples for which the target entity was ranked in the top 10 predictions), which were originally used in the entity prediction task (Bordes et al., 2013), recent work also reports the mean reciprocal rank (**MRR**) which is commonly used in information retrieval. In both "Raw" and "Filtered" settings, mean rank is always greater or equal to 1 and the lower mean rank indicates better entity prediction performance. MRR and Hits@10 scores always range from 0.0 to 1.0, and higher score reflects better prediction result.

3.1.2 Main results

Table 3 lists the entity prediction results of KB completion models on the WN18 and FB15k datasets. The first 18 rows report the performance of the models that do not exploit information about alternative paths between head and tail entities. The next 5 rows report results of the models that exploit information about relation paths. The last 3 rows present results for the models which make use of textual mentions derived from a large external corpus. It is clear that the models with the additional external corpus information obtained the best results. Table 3 also shows that the models employing path information generally achieve better results than models that do not use such information. In terms of models not exploiting path information or external information, the ComplEx model obtains highest scores on FB15k. In addition, the STransE model performs better than its closely related models SE, TransE, TransR, CTransR, TransD and TranSparse on both WN18 and FB15k.

3.2 Triple classification

3.2.1 Task description

The triple classification task was first introduced by Socher et al. (2013), and since then it has been used to evaluate various embedding models. The aim of this task is to predict whether a triple (h, r, t) is correct or not. For classification, a relation-specific threshold θ_r is set for each relation type r. If the implausibility score of an unseen test triple (h, r, t) is smaller than θ_r then the triple will be classified as correct, otherwise incorrect. Following Socher et al. (2013), the relation-specific thresholds are determined by maximizing the micro-averaged accuracy, which is a per-triple average, on the validation set.

3.2.2 Main results

Table 4 presents the triple classification results of KB completion models on the WN11 and FB13 datasets.[1] The first 6 rows report the performance of models that use TransE to initialize the entity and relation vectors. The last 11 rows present the accuracy of models with randomly initialized parameters. Note that there are higher results reported for NTN, Bilinear-COMP and TransE-COMP when entity vectors are initialized by averaging the pre-trained word vectors (Mikolov et al., 2013; Pennington et al., 2014). It is not surprising as many entity names in WordNet and FreeBase are lexically meaningful. It is possible for all other embedding models to utilize the pre-trained word vectors as well. However, as pointed out by Wang et al. (2014) and Guu et al. (2015), averaging the pre-trained word vectors for initializing entity vectors is an open problem and it is not always useful since entity names in many domain-specific KBs are not lexically meaningful.

Method	**W11**	**F13**	Avg.
TransR (Lin et al., 2015b)	85.9	82.5	84.2
CTransR (Lin et al., 2015b)	85.7	-	-
TransD (Ji et al., 2015)	86.4	**89.1**	**87.8**
TEKE_H (Wang and Li, 2016)	84.8	84.2	84.5
TranSparse-S (Ji et al., 2016)	86.4	88.2	87.3
TranSparse-US (Ji et al., 2016)	86.8	87.5	87.2
NTN (Socher et al., 2013)	70.6	87.2	78.9
TransH (Wang et al., 2014)	78.8	83.3	81.1
SLogAn (Liang and Forbus, 2015)	75.3	85.3	80.3
KG2E (He et al., 2015)	85.4	85.3	85.4
Bilinear-COMP (Guu et al., 2015)	77.6	86.1	81.9
TransE-COMP (Guu et al., 2015)	80.3	87.6	84.0
TransR-FT (Feng et al., 2016a)	86.6	82.9	84.8
TransG (Xiao et al., 2016)	**87.4**	87.3	87.4
lppTransD (Yoon et al., 2016)	86.2	88.6	87.4
TransE (Bordes et al., 2013) [*]	85.2	87.6	86.4
TransE-NMM (Nguyen et al., 2016a)	86.8	88.6	87.7

Table 4: Accuracy results (in %) for triple classification on WN11 (labeled as **W11**) and FB13 (labeled as **F13**) test sets. "Avg." denotes the averaged accuracy. [*]: TransE results from the implementation of Nguyen et al. (2016a) because these results are higher than those first reported in Wang et al. (2014).

4 Conclusions and further discussion

This article presented an overview of embedding models of entity and relationships for KB completion. The article also provided update-to-date experimental results of the embedding models on the entity prediction and triple classification tasks.

It would be interesting to further explore those models for a new application where triple-style data is available. For example, Vu et al. (2017) extended the STransE model (Nguyen et al., 2016b) for the *search personalization* task in information retrieval, to model a *user*-oriented relationship between *query* and *document*.

[1] In fact, as shown in tables 3 and 4, TransE obtains a very competitive performance on entity prediction and triple classification tasks. The similar observation is also found by García-Durán et al. (2015), Lin et al. (2015a) and García-Durán et al. (2016). The reason is due to a careful grid search.

References

[Bollacker et al.2008] Kurt Bollacker, Colin Evans, Praveen Paritosh, Tim Sturge, and Jamie Taylor. 2008. Freebase: A Collaboratively Created Graph Database for Structuring Human Knowledge. In *Proceedings of the 2008 ACM SIGMOD International Conference on Management of Data*, pages 1247–1250.

[Bordes et al.2011] Antoine Bordes, Jason Weston, Ronan Collobert, and Yoshua Bengio. 2011. Learning Structured Embeddings of Knowledge Bases. In *Proceedings of the Twenty-Fifth AAAI Conference on Artificial Intelligence*, pages 301–306.

[Bordes et al.2012] Antoine Bordes, Xavier Glorot, Jason Weston, and Yoshua Bengio. 2012. A Semantic Matching Energy Function for Learning with Multi-relational Data. *Machine Learning*, 94(2):233–259.

[Bordes et al.2013] Antoine Bordes, Nicolas Usunier, Alberto Garcia-Duran, Jason Weston, and Oksana Yakhnenko. 2013. Translating Embeddings for Modeling Multi-relational Data. In *Advances in Neural Information Processing Systems 26*, pages 2787–2795.

[Dong et al.2014] Xin Dong, Evgeniy Gabrilovich, Geremy Heitz, Wilko Horn, Ni Lao, Kevin Murphy, Thomas Strohmann, Shaohua Sun, and Wei Zhang. 2014. Knowledge vault: A web-scale approach to probabilistic knowledge fusion. In *Proceedings of the 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 601–610.

[Duchi et al.2011] John Duchi, Elad Hazan, and Yoram Singer. 2011. Adaptive Subgradient Methods for Online Learning and Stochastic Optimization. *The Journal of Machine Learning Research*, 12:2121–2159.

[Feng et al.2016a] Jun Feng, Minlie Huang, Mingdong Wang, Mantong Zhou, Yu Hao, and Xiaoyan Zhu. 2016a. Knowledge graph embedding by flexible translation. In *Principles of Knowledge Representation and Reasoning: Proceedings of the Fifteenth International Conference*, pages 557–560.

[Feng et al.2016b] Jun Feng, Minlie Huang, Yang Yang, and xiaoyan zhu. 2016b. GAKE: Graph Aware Knowledge Embedding. In *Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: Technical Papers*, pages 641–651.

[García-Durán et al.2015] Alberto García-Durán, Antoine Bordes, and Nicolas Usunier. 2015. Composing Relationships with Translations. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 286–290.

[García-Durán et al.2016] Alberto García-Durán, Antoine Bordes, Nicolas Usunier, and Yves Grandvalet. 2016. Combining Two and Three-Way Embedding Models for Link Prediction in Knowledge Bases. *Journal of Artificial Intelligence Research*, 55:715–742.

[Gardner and Mitchell2015] Matt Gardner and Tom Mitchell. 2015. Efficient and Expressive Knowledge Base Completion Using Subgraph Feature Extraction. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 1488–1498.

[Gardner et al.2014] Matt Gardner, Partha P. Talukdar, Jayant Krishnamurthy, and Tom M. Mitchell. 2014. Incorporating Vector Space Similarity in Random Walk Inference over Knowledge Bases. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing*, pages 397–406.

[Guu et al.2015] Kelvin Guu, John Miller, and Percy Liang. 2015. Traversing Knowledge Graphs in Vector Space. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 318–327.

[He et al.2015] Shizhu He, Kang Liu, Guoliang Ji, and Jun Zhao. 2015. Learning to Represent Knowledge Graphs with Gaussian Embedding. In *Proceedings of the 24th ACM International on Conference on Information and Knowledge Management*, pages 623–632.

[Jenatton et al.2012] Rodolphe Jenatton, Nicolas L. Roux, Antoine Bordes, and Guillaume R Obozinski. 2012. A latent factor model for highly multi-relational data. In *Advances in Neural Information Processing Systems 25*, pages 3167–3175.

[Ji et al.2015] Guoliang Ji, Shizhu He, Liheng Xu, Kang Liu, and Jun Zhao. 2015. Knowledge Graph Embedding via Dynamic Mapping Matrix. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 687–696.

[Ji et al.2016] Guoliang Ji, Kang Liu, Shizhu He, and Jun Zhao. 2016. Knowledge Graph Completion with Adaptive Sparse Transfer Matrix. In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence*, pages 985–991.

[Krompaß et al.2015] Denis Krompaß, Stephan Baier, and Volker Tresp. 2015. Type-Constrained Representation Learning in Knowledge Graphs. In *Proceedings of the 14th International Semantic Web Conference*, pages 640–655.

[Lao and Cohen2010] Ni Lao and William W. Cohen. 2010. Relational retrieval using a combination of path-constrained random walks. *Machine Learning*, 81(1):53–67.

[Lao et al.2011] Ni Lao, Tom Mitchell, and William W. Cohen. 2011. Random Walk Inference and Learning in a Large Scale Knowledge Base. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, pages 529–539.

[Lehmann et al.2015] Jens Lehmann, Robert Isele, Max Jakob, Anja Jentzsch, Dimitris Kontokostas, Pablo N. Mendes, Sebastian Hellmann, Mohamed Morsey, Patrick van Kleef, Sören Auer, and Christian Bizer. 2015. DBpedia - A Large-scale, Multilingual Knowledge Base Extracted from Wikipedia. *Semantic Web*, 6(2):167–195.

[Liang and Forbus2015] Chen Liang and Kenneth D. Forbus. 2015. Learning Plausible Inferences from Semantic Web Knowledge by Combining Analogical Generalization with Structured Logistic Regression. In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence*, pages 551–557.

[Lin et al.2015a] Yankai Lin, Zhiyuan Liu, Huanbo Luan, Maosong Sun, Siwei Rao, and Song Liu. 2015a. Modeling Relation Paths for Representation Learning of Knowledge Bases. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 705–714.

[Lin et al.2015b] Yankai Lin, Zhiyuan Liu, Maosong Sun, Yang Liu, and Xuan Zhu. 2015b. Learning Entity and Relation Embeddings for Knowledge Graph Completion. In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence Learning*, pages 2181–2187.

[Liu and Nocedal1989] D. C. Liu and J. Nocedal. 1989. On the Limited Memory BFGS Method for Large Scale Optimization. *Mathematical Programming*, 45(3):503–528.

[Liu et al.2016] Qiao Liu, Liuyi Jiang, Minghao Han, Yao Liu, and Zhiguang Qin. 2016. Hierarchical Random Walk Inference in Knowledge Graphs. In *Proceedings of the 39th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 445–454.

[Luo et al.2015] Yuanfei Luo, Quan Wang, Bin Wang, and Li Guo. 2015. Context-Dependent Knowledge Graph Embedding. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 1656–1661.

[Mikolov et al.2013] Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. 2013. Distributed Representations of Words and Phrases and their Compositionality. In *Advances in Neural Information Processing Systems 26*, pages 3111–3119.

[Miller1995] George A. Miller. 1995. WordNet: A Lexical Database for English. *Communications of the ACM*, 38(11):39–41.

[Neelakantan et al.2015] Arvind Neelakantan, Benjamin Roth, and Andrew McCallum. 2015. Compositional Vector Space Models for Knowledge Base Completion. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 156–166.

[Nguyen et al.2016a] Dat Quoc Nguyen, Kairit Sirts, Lizhen Qu, and Mark Johnson. 2016a. Neighborhood Mixture Model for Knowledge Base Completion. In *Proceedings of The 20th SIGNLL Conference on Computational Natural Language Learning*, pages 40–50.

[Nguyen et al.2016b] Dat Quoc Nguyen, Kairit Sirts, Lizhen Qu, and Mark Johnson. 2016b. STransE: a novel embedding model of entities and relationships in knowledge bases. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 460–466.

[Nickel et al.2011] Maximilian Nickel, Volker Tresp, and Hans-Peter Kriegel. 2011. A Three-Way Model for Collective Learning on Multi-Relational Data. In *Proceedings of the 28th International Conference on Machine Learning*, pages 809–816.

[Nickel et al.2016a] Maximilian Nickel, Kevin Murphy, Volker Tresp, and Evgeniy Gabrilovich. 2016a. A Review of Relational Machine Learning for Knowledge Graphs. *Proceedings of the IEEE*, 104(1):11–33.

[Nickel et al.2016b] Maximilian Nickel, Lorenzo Rosasco, and Tomaso Poggio. 2016b. Holographic embeddings of knowledge graphs. In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence*, pages 1955–1961.

[Niepert2016] Mathias Niepert. 2016. Discriminative Gaifman Models. In *Advances in Neural Information Processing Systems 29*, pages 3405–3413.

[Pennington et al.2014] Jeffrey Pennington, Richard Socher, and Christopher Manning. 2014. Glove: Global Vectors for Word Representation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing*, pages 1532–1543.

[Shi and Weninger2017] Baoxu Shi and Tim Weninger. 2017. ProjE: Embedding Projection for Knowledge Graph Completion. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence*.

[Socher et al.2013] Richard Socher, Danqi Chen, Christopher D Manning, and Andrew Ng. 2013. Reasoning With Neural Tensor Networks for Knowledge Base Completion. In *Advances in Neural Information Processing Systems 26*, pages 926–934.

[Suchanek et al.2007] Fabian M. Suchanek, Gjergji Kasneci, and Gerhard Weikum. 2007. YAGO: A Core of Semantic Knowledge. In *Proceedings of the 16th International Conference on World Wide Web*, pages 697–706.

[Taskar et al.2004] Ben Taskar, Ming fai Wong, Pieter Abbeel, and Daphne Koller. 2004. Link Prediction in Relational Data. In *Advances in Neural Information Processing Systems 16*, pages 659–666.

[Tay et al.2017] Yi Tay, Anh Tuan Luu, Siu Cheung Hui, and Falk Brauer. 2017. Random Semantic Tensor Ensemble for Scalable Knowledge Graph Link Prediction. In *Proceedings of the Tenth ACM International Conference on Web Search and Data Mining*, pages 751–760.

[Toutanova and Chen2015] Kristina Toutanova and Danqi Chen. 2015. Observed Versus Latent Features for Knowledge Base and Text Inference. In *Proceedings of the 3rd Workshop on Continuous Vector Space Models and their Compositionality*, pages 57–66.

[Toutanova et al.2016] Kristina Toutanova, Victoria Lin, Wen-tau Yih, Hoifung Poon, and Chris Quirk. 2016. Compositional Learning of Embeddings for Relation Paths in Knowledge Base and Text. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1434–1444.

[Trouillon et al.2016] Théo Trouillon, Johannes Welbl, Sebastian Riedel, Éric Gaussier, and Guillaume Bouchard. 2016. Complex Embeddings for Simple Link Prediction. In *Proceedings of the 33nd International Conference on Machine Learning*, pages 2071–2080.

[Vu et al.2017] Thanh Vu, Dat Quoc Nguyen, Mark Johnson, Dawei Song, and Alistair Willis. 2017. Search Personalization with Embeddings. In *Proceedings of the 39th European Conference on Information Retrieval*.

[Wang and Li2016] Zhigang Wang and Juan-Zi Li. 2016. Text-Enhanced Representation Learning for Knowledge Graph. In *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence*, pages 1293–1299.

[Wang et al.2014] Zhen Wang, Jianwen Zhang, Jianlin Feng, and Zheng Chen. 2014. Knowledge Graph Embedding by Translating on Hyperplanes. In *Proceedings of the Twenty-Eighth AAAI Conference on Artificial Intelligence*, pages 1112–1119.

[Wang et al.2016] Quan Wang, Jing Liu, Yuanfei Luo, Bin Wang, and Chin-Yew Lin. 2016. Knowledge Base Completion via Coupled Path Ranking. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1308–1318.

[Wei et al.2016] Zhuoyu Wei, Jun Zhao, and Kang Liu. 2016. Mining Inference Formulas by Goal-Directed Random Walks. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1379–1388.

[West et al.2014] Robert West, Evgeniy Gabrilovich, Kevin Murphy, Shaohua Sun, Rahul Gupta, and Dekang Lin. 2014. Knowledge Base Completion via Search-based Question Answering. In *Proceedings of the 23rd International Conference on World Wide Web*, pages 515–526.

[Xiao et al.2016] Han Xiao, Minlie Huang, and Xiaoyan Zhu. 2016. TransG : A Generative Model for Knowledge Graph Embedding. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2316–2325.

[Xiao et al.2017] Han Xiao, Minlie Huang, and Xiaoyan Zhu. 2017. SSP: semantic space projection for knowledge graph embedding with text descriptions. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence*.

[Yang et al.2015] Bishan Yang, Wen-tau Yih, Xiaodong He, Jianfeng Gao, and Li Deng. 2015. Embedding Entities and Relations for Learning and Inference in Knowledge Bases. In *Proceedings of the International Conference on Learning Representations*.

[Yoon et al.2016] Hee-Geun Yoon, Hyun-Je Song, Seong-Bae Park, and Se-Young Park. 2016. A Translation-Based Knowledge Graph Embedding Preserving Logical Property of Relations. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 907–916.

[Zeiler2012] Matthew D. Zeiler. 2012. ADADELTA: An Adaptive Learning Rate Method. *CoRR*, abs/1212.5701.